

Open Community Runtime: A Framework for Cooperative Resource Control in Exascale Systems

Vivek Sarkar, vsarkar@rice.edu [Contact author]
Barbara Chapman, chapman@cs.uh.edu
Bill Gropp, wgropp@uiuc.edu
Rob Knauerhase, rob.knauerhase@intel.com
Tim Mattson, timothy.g.mattson@intel.com
Wilf Pinfold, wilfred.pinfold@intel.com

Challenges Addressed: Exascale systems will be qualitatively different from current and past HPC systems. Specifically, they will use massive numbers of multi-core processors with hundreds of homogeneous and heterogeneous cores per chip, their performance will be driven by parallelism and constrained by energy and data movement. They will also be subject to frequent faults and failures. Unlike previous generations of hardware evolution, these trends will have a profound impact on the foundations of future HPC software [5].

The design of Operating Systems and Runtimes for past HPC systems assumed a layered software architecture. Specific layers were isolated through well-defined interfaces allowing innovation within each layer to proceed independent of the other layers. The development of MPI, for example, enabled advances in the communication runtime layer to proceed independently of advances of thread scheduling and other forms of resource control in the operating system layer. Our position is that this form of layered architecture will not be effective for exascale Systems. A new approach based on cooperative resource control across multiple OS/R layers will be essential to address the programmability, concurrency, energy, locality, and resilience challenges faced by exascale systems. In current approaches, each software layer has its own policies and mechanisms for managing resources such as processors/threads, memories, and network interfaces using abstract models that are oblivious to the presence of other layers sharing the same or related resources. The limitations of such approaches are manifest, for example, in the difficulties faced by attempts to combine scheduling and synchronization of MPI processes with scheduling and synchronization of intra-process threads and tasks. In contrast, effective solutions to cooperative resource control will enable combinations of OS/R components (communication, scheduling, synchronization, memory management, and resilience) to work together in ensuring that key resource constraints are satisfied *in a holistic manner*. Finally, the runtime system must play a central role in addressing programmability challenges by bridging the wide semantic gap between high-level programming models and the complexity of exascale hardware.

Novelty: We propose a research agenda comprehending the development of a shared framework for cooperative resource control that we call the Open Community Runtime (OCR). Extreme Scale computing is limited by concurrency, energy efficiency, and resiliency. The runtime system will play a critical role in exascale enablement for multiple reasons because the inherent variability in extreme scale software and hardware components calls for end-to-end asynchrony in system design, and tight integration of inter-node and intra-node parallel runtime systems has proven elusive thus far. The term “runtime” includes lightweight OS services that can be made available in user mode. Further, we expect that the lessons learned on runtime systems from this proposed research will also influence kernel design in future operating systems [14].

The framework will enable cooperative resource control by including first-class support for negotiation of resources across runtime services, and will allow for multiple implementations of common runtime APIs, perhaps with different trade-offs in different dimensions of resource usage. An exemplar of the OCR framework will be made available in open source form. Software projects that the authors of this position paper and their collaborators are involved with will serve as the starting point for development of OCR.

OCR will support innovations in programmability by providing a common substrate for researchers exploring a wide range of programming models. The authors of this position paper are interested in exploring a range of programming models on OCR that include Concurrent Collections [13,15], Habanero-C [13], MPI, and OpenMP. We will emphasize modern programming models (including future versions of MPI and OpenMP) that generate small units of work for computation and communication mapped by OCR onto coarser-grained system resources such as processors and network interfaces, while meditating the control of resources to ensure global constraints for resource usage and resilience are satisfied. Furthermore, since OCR will map onto a range of extremely scalable systems, it will provide a common software infrastructure underneath which hardware researchers can freely innovate.

There are many differences between our proposed approach and existing solutions. Past runtime systems

have either been optimized for dynamic parallelism that is oblivious of locality (e.g., Cilk, OpenMP, Intel Thread Building Blocks) or for locality in the absence of dynamic parallelism (e.g., MPI, shmem, UPC). In contrast, OCR will support unbounded amounts of dynamic parallelism [2] with locality control [7] so as to be "forward scalable" to multiple generations of parallel hardware. MPI and PGAS programming models only express two levels of locality -- local and remote. HPCS languages such as Chapel and X10 mitigate this by allowing the programmer to express dynamic parallelism with "locales" and "places", but their locality model is also restricted to two levels. In contrast, OCR will support hierarchical place abstractions with arbitrary depths [4] that include support for heterogeneous accelerators [10,13]. Synchronization goes hand-in-hand with communication. Classical approaches to intra-node synchronization involved blocking of tasks at key events such as collective operations (e.g. reductions). In contrast, OCR will make these events as asynchronous as possible through the support of primitives such as codelets [9], data-driven tasks [11], phasers [1,3,6,8], and delegated isolation [12]. In summary, we are unaware of any existing solution that focuses on an open source runtime framework like the OCR while supporting a wide range of choices in programming models and system hardware.

Uniqueness: Components of OCR will be useful for sub-exascale systems, hence at the level of individual components, there is considerable overlap between OCR and current systems. We are unaware, however, of other research programs that could realize the integration across the runtime solution stack for exascale systems as planned for OCR. While the urgency for OCR is perhaps lower on sub-exascale systems since a single programming model (e.g., MPI+X) may suffice and energy constraints are less severe, it is critical to make OCR available to application developers soon so they can prepare their codes for the exascale era.

The resilience problem, i.e. the requirement that a system will continue to make progress in the face of failures of individual components, is less critical for sub-exascale systems. OCR will differ from established practice in current systems by directly addressing the resilience problems with fundamental research into power and performance efficient APIs that are resilient and enable a dynamic negotiation of resource sharing, as well as introspective self-observation to detect/anticipate and adapt to faults.

Applicability: While the emphasis is on exascale systems, components of OCR will be applicable to extreme scale systems at sub-exascale levels (e.g., departmental petascale, embedded terascale). Hence, hardware researchers (such as the co-authors on this proposal from Intel) responding to challenges raised by OCR could develop hardware innovations that would impact mainstream microprocessors. Likewise, the programming models enabled by OCR could also be helpful in addressing the very-high-core-count software challenges facing the computer industry.

Maturity and Effort: The authors of this proposal have served as key participants in exascale workshops, studies, and reports through which they have gained a shared understanding of the OS/R challenges facing exascale systems. This shared understanding has deepened through the lessons learned in the UHPC Runnemedede project led by Intel. Further, the authors plan to leverage existing software developed in their groups or by their collaborators as a starting point for development of OCR (Habanero-C, Intel codelets, MPICH-2 Nemesis, OpenUH runtime, qthreads, X10 distributed runtime). At the same time, the absence of any backward compatibility requirements for OCR ensures that it can focus on providing clear, exascale-oriented APIs enabling innovation without any legacy commitments. Thereby, it can provide flexibility in design, implementation, and schedule while its key results can remain applicable/portable to legacy programming models. We envision a core effort that is under 10PY/year for 3 years to realize the proposed OCR framework, with annual intermediate releases. The ultimate goal is for the effort to grow on a community-wide basis as researchers beyond the core team use the OCR for hardware and software exploration and also contribute additional components to OCR.

References

- [1] Phasers: a Unified Deadlock-Free Construct for Collective and Point-to-point Synchronization. Jun Shirako, David Peixotto, Vivek Sarkar, William Scherer. Proceedings of the 2008 ACM International Conference on Supercomputing (ICS), June 2008.
- [2] Work-First and Help-First Scheduling Policies for Terminally Strict Parallel Programs. Yi Guo, Rajkishore Barik, Raghavan Raman, Vivek Sarkar. 23rd IEEE International Parallel and Distributed Processing Symposium (IPDPS), May 2009.
- [3] Phaser Accumulators: a New Reduction Construct for Dynamic Parallelism. Jun Shirako, David Peixotto, Vivek Sarkar, William Scherer. 23rd IEEE International Parallel and Distributed Processing Symposium (IPDPS), May 2009
- [4] Hierarchical Place Trees: A Portable Abstraction for Task Parallelism and Date Movement. Yonghong Yan, Jisheng Zhao, Yi Guo, Vivek Sarkar. Proceedings of the 22nd Workshop on Languages and Compilers for Parallel Computing (LCPC), October 2009.
- [5] Software Challenges in Extreme Scale Systems. V. Sarkar, W. Harrod, A.E. Snively. SciDAC Review Special Issue on Advanced Computing: The Roadmap to Exascale, pp. 60-65, January 2010.
- [6] Hierarchical Phasers for Scalable Synchronization and Reduction. Jun Shirako, Vivek Sarkar. 24th IEEE International Parallel and Distributed Processing Symposium (IPDPS), April 2010.
- [7] SLAW: a Scalable Locality-aware Adaptive Work-stealing Scheduler. Yi Guo, Jisheng Zhao, Vincent Cavé, Vivek Sarkar. 24th IEEE International Parallel and Distributed Processing Symposium (IPDPS), April 2010.
- [8] Unifying Barrier and Point-to-Point Synchronization in OpenMP with Phasers. Jun Shirako, Kamal Sharma, Vivek Sarkar. 7th International Workshop on OpenMP (IWOMP), June 2011.
- [9] Using a Codelet Program Execution Model for Exascale Machines. Stephane Zuckerman, Joshua Suetterlein, Rob Knauerhase, Guang Gao. 1st International Workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era (EXADAPT '11), June 2011.
- [10] Dynamic Task Parallelism with a GPU Work-Stealing Runtime System. Sanjay Chatterjee, Max Grossman, Alina Sbirlea, Vivek Sarkar. 2011 Workshop on Languages and Compilers for Parallel Computing (LCPC), September 2011.
- [11] Data-Driven Tasks and their Implementation. Sagnak Tasirlar, Vivek Sarkar. Proceedings of the International Conference on Parallel Processing (ICPP) 2011, September 2011.
- [12] Delegated Isolation. Roberto Lublinerman, Jisheng Zhao, Zoran Budimlic, Swarat Chaudhuri, Vivek Sarkar. Proceedings of OOPSLA 2011, October 2011.
- [13] Mapping a Data-Flow Programming Model onto Heterogeneous Platforms. Alina Sbirlea, Yi Zou, Zoran Budimlic, Jason Cong, Vivek Sarkar. Conference on Languages, Compilers, Tools and Theory for Embedded Systems (LCTES), June 2012.
- [14] For Extreme Parallelism, Your OS is Soooo Last Millennium". Rob Knauerhase, Romain Cledat, Justin Teller. USENIX Hot Topics in Parallelism (HotPar), June 2012.
- [15] Concurrent Collections on Distributed Memory Theory Put Into Practice. F.Schlimbach, J.Brodman, K. Knobe. International European Conference on Parallel and Distributed Computing (Euro-Par), August 2012.